



WEST Search History

[Hide Items](#) [Restore](#) [Clear](#) [Cancel](#)

DATE: Wednesday, October 05, 2005

<u>Hide?</u>	<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>			
<input type="checkbox"/>	L13	L12 and (shared bandwidth)	0
<input type="checkbox"/>	L12	L11 and (priority or prioritize or prioritization)	31
<input type="checkbox"/>	L11	20000417	89
<input type="checkbox"/>	L10	L9 and (queue or queuing)	279
<input type="checkbox"/>	L9	TCP socket	613
<input type="checkbox"/>	L8	20000417	7
<input type="checkbox"/>	L7	(manage or management or managing or control or controlling or manipulate or manipulating) near8 (shared bandwidth)	41
<input type="checkbox"/>	L6	L5 and socket	0
<input type="checkbox"/>	L5	20000417	2
<input type="checkbox"/>	L4	front-end near8 (interface or interfacing) near8 (priority or prioritize or prioritization)	3
<input type="checkbox"/>	L3	L2 and (priority or prioritize or prioritization)	9
<input type="checkbox"/>	L2	20000417	29
<input type="checkbox"/>	L1	(last-mile or (last mile)) near8 (communication or connectivity)	202

END OF SEARCH HISTORY

WEST Search History

DATE: Wednesday, October 05, 2005

<u>Hide?</u>	<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>			
<input type="checkbox"/>	L16	l15 and priority	0
<input type="checkbox"/>	L15	20000417	3
<input type="checkbox"/>	L14	(front-end or (front end)) adj manager	28
<input type="checkbox"/>	L13	L12 and (prioritize or prioritizing or prioritized)	23
<input type="checkbox"/>	L12	20000417	212
<input type="checkbox"/>	L11	TCP adj (channel or socket)	678
<input type="checkbox"/>	L10	L8 and l4	0
<input type="checkbox"/>	L9	L8 and l5	0
<input type="checkbox"/>	L8	TCP adj2 (channel or socket)	2184
<input type="checkbox"/>	L7	l5 and TCP	3
<input type="checkbox"/>	L6	lr and TCP	1223
<input type="checkbox"/>	L5	L4 and (priority adj2 (level or value))	24
<input type="checkbox"/>	L4	L3 and l2	78
<input type="checkbox"/>	L3	(channel or socket) near5 buffer	20356
<input type="checkbox"/>	L2	20000417	493
<input type="checkbox"/>	L1	(prioritize or prioritizing or prioritized) near8 (socket or channel)	932

END OF SEARCH HISTORY

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L3: Entry 9 of 9

File: USPT

May 26, 1998

DOCUMENT-IDENTIFIER: US 5757784 A

TITLE: Usage-based billing system for full mesh multimedia satellite network

Application Filing Date (1):

19960104

Brief Summary Text (4):

Recent technological advances have made satellite technology a viable option for inter-office communications. Satellite communication systems provide several advantages over terrestrial communications systems for networking applications. Terrestrial based communication systems are typically subject to the "last mile" problem, wherein circuits providing increased bandwidth for terrestrial communication are not compatible with existing, and typically analog, circuits available at most commercial facilities. Satellite communication avoids this problem by allowing transmission directly to and from the roof tops, for example, of commercial facilities within a network. Satellite communication systems are characterized by distance insensitivity within the coverage area of the satellite and are not subject to interference as are terrestrial communication systems. Satellite communication systems also have an inherent broadcast capability which simplifies routing in a full mesh network and facilitates multi-party videoconferencing. These benefits make satellite communications a preferred choice for multi-national, multi-media networks, especially where some or all of the sites are located remotely from a reliable communications infrastructure.

Brief Summary Text (8):

In accordance with an embodiment of the present invention, a satellite communication system is provided which uses TDMA/DAMA (time division multiple access/demand assigned multiple access) switching technology to efficiently transmit voice, video and bursty data. TDMA/DAMA switching provides significant cost benefits due to shared use of the space segment. Demanding real time transmissions, such as voice and video data, share bandwidth with bursty data streams from local area networks. Users can order low bit rate services for most of the day, and acquire and pay for higher rate channels only when they are needed. Video channels can be acquired on-demand, and paid for only as long as they are used. Voice calls are connected as needed. Since voice is digitized and packetized, silence periods can be used for the transmission of lower priority data file transfers. The system of the present invention provides services for data, and dial-up services for voice and video. Multi-user video conferencing can also be provided.

Detailed Description Text (8):

If there are insufficient fragments to fill a burst buffer 68, the buffer 68 is padded with null characters. If there are more fragments available than can be stored in a burst buffer 68, a priority scheme is used to determine which fragments are stored in the buffer 68. In accordance with an embodiment of the present invention, the priority scheme preferably provides two primary levels of priority, and a number of sublevels of priority. The first primary priority level is for real time data. The second primary priority level is for non-real time data. The classification of real and non-real time is arbitrary, and can be assigned using

the NMC 13. Real time data is usually time-dependent data such as voice and video. Priority can be assigned by the port, or by the address of a virtual circuit. Non-real time data priority levels can be further determined by assigning CIRs to the port or virtual circuit.

Detailed Description Text (14):

In summary, data frames directed towards a terminal 12 by local user access devices 43 are fragmented for transmission efficiency, multiplexed according to a priority scheme, and stored in a burst buffer 68 for burst transmission at a prescribed time. TDMA bursts are transmitted to a satellite 14, where they are broadcast on a downlink to all other terminals 12 tuned to the carrier.

Detailed Description Text (19):

The FAD 66 assembles the fragments in the payloads 108 of bursts retained by the terminal 12 into the frames from which the fragments were derived using fragmentation overhead bytes. The fragmentation overhead bytes are described in further detail below in connection with FIG. 7. Incomplete frames are discarded, and an error report is generated by the FAD 66 for use by the NMC 13. Each complete frame contains an address, whose form is dependent upon the nature of the frame, as stated previously. In accordance with a second level of filtering, the address is analyzed to determine if it is associated with one of the communications ports 40 attached to the terminal 12. If so, the frame is forwarded by the FAD 66 to the appropriate Frame Handler module (FHM) 64. If the address is not associated with one of the communication ports 40 or the NMM 78 (FIG. 4), the frame is discarded. The FHM 64, using priority queuing, transmits the frame to the appropriate communications port 40. For example, a frame containing time critical information such as voice or video is transmitted to the appropriate port 40 before a frame containing, for example, data that is not time critical. The FHM 64 records statistics regarding the volume of data transmitted by a virtual circuit and a port which are also used by the NMC 13.

Detailed Description Text (32):

As stated previously, the fragmentation overhead bytes 118 uniquely associate a fragment to the frame from which it was derived, and identify where in the burst payload 106 the fragment starts. For example, if a payload is 300 bytes, and fragments are each 64 bytes, the fragmentation overhead bytes can comprise a byte for each fragment having at least 12 bits to identify where in the payload the fragment begins, and another set of bits to specify a unique identification number for the frame from which the fragment was derived. The frame identification number need not be an address in the frame, but rather a number arbitrarily assigned to each frame by the FAD 66, for example, for all of the bursts 98 transmitted by the terminal 12 during a TDMA frame 96. The payload 108 of a burst typically comprises both Ethernet and frame relay formatted data. It may also contain other types of frames such as HDLC and SDLC. Fragmentation is a generic format in that it allows these different types of data frames to be transmitted in the same payload 108. Further, frames in the payload 108 are fragmented to ensure that the burst payload 108 can be efficiently packed (i.e., the payload contains as few null characters as possible), and to prevent oversized fragments from delaying transmission of higher priority fragments such as voice. The fragment size is a system parameter, and is typically set at 64 or 128 bytes.

Detailed Description Text (51):

The system 10 is advantageous because it employs statistical multiplexing within each terminal 12 and between terminals 12 to provide cost effective sharing of the satellite link 16. The burst buffer 68 in each terminal 12 allows bandwidth sharing within a terminal 12. All data streams share the burst buffer 68 regardless of whether they contain voice, video or data. As stated previously, the system 10 uses a priority scheme to fill the burst buffer 68 to capacity as much as possible. Statistical multiplexing occurs within a terminal 12 because not all permanent virtual circuits (PVCs) and voice calls are constantly busy. Thus, the system 10

uses silence intervals in voice conversations, which are detected by the voice multiplexer 62, to transmit lower priority data packets, thereby increasing the effective bandwidth available for voice. The system 10 also takes advantage of the bursty nature of data in a multi-LAN environment. A terminal 12 can multiplex data from one Ethernet LAN and a number of router-based LANs. The data frames can, in turn, be statistically multiplexed with voice frames in the burst buffer 68.

Detailed Description Text (54):

A terminal 12 with relatively low bandwidth demand can release its bandwidth for use by other terminals 12 requiring more capacity. If standard CIR service is used, even the bandwidth associated with a constant bit rate (CBR) service can be released to the common pool maintained by the terminal 12 operating as the MT if the CBR service terminal 12 is not using the bandwidth. In such cases, a terminal 12 with lower priority, but momentarily demanding throughput needs, can borrow the bandwidth from the common pool. The borrowed bandwidth is subsequently returned to the lending terminal 12 when needed. Thus, a user can achieve higher burst rates without having to pay for a CBR service equal to the desired burst rate.

Detailed Description Text (58):

Voice and video can be transmitted over the satellite link 16 via a terminal 12. Unlike known frame relay systems, the system 10 accommodates voice within the CIR. The user is billed a fixed price for the CIR. The system 10 does not assign bandwidth outside of the terminal CIR to accommodate telephone calls. In all cases generally, a voice data stream is given priority over data, an advantage not available in known frame relay systems which use voice multiplexers connected to the frame relay and do not give voice data streams priority.

Detailed Description Text (71):

An SCPC system such as that disclosed in U.S. Pat. No. 5,434,850, to Fielding et al, uses frame relay to broadcast data and voice over a satellite link. The system 10 disclosed herein uses a cell-based protocol which allows more effective transmission of time critical information such as voice and video. The system disclosed in Fielding et al. patent supports a four level priority scheme based on DLCI. The system 10 supports more levels of priority in the data area, by means of its ability to support CIRs for the virtual circuits. The Fielding system does not support CIRs.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L3: Entry 4 of 9

File: USPT

Mar 12, 2002

DOCUMENT-IDENTIFIER: US 6356537 B1

TITLE: Radio interface card for a broadband wireless ATM system

Application Filing Date (1):

19980924

Brief Summary Text (4):

Broadband wireless systems such as Local Multipoint Distribution Systems (LMDS), known as Local Multipoint Communication System (LMCS) in Canada, are being developed to provide point to multipoint, high bandwidth services between a base station connected to a backbone such as an asynchronous transfer mode (ATM) network and network interface units (NIUs) at fixed locations within a defined geographic area or cell. A wireless link between the base station and the NIUs operates at a wireless radio frequency (RF) typically in the 2.5 to 40 GHz range depending on the allocated frequency license. A transmitter/receiver at the base station and a transceiver at each NIU site supports bi-directional, broadband "last mile" communication between a service provider and a customer.

Detailed Description Text (16):

1. Backplane Adaptation 60: Data is either segmented (for an Adaptation Card) or mapped (for a Cell Relay Card) into ATM-like Cells before being transmitted on the Local Add Bus to the Hub Card. Seven bytes of Newbridge header overhead are added to the five bytes of ATM overhead to form a 60 byte Newbridge Cell. The Newbridge header contains the Priority of the attached ATM Cell, the UCS Destination Address as well as whether the Cell is a point-to-point cell or a Multicast Cell.

Detailed Description Text (18):

2. Arbitration Queuing 62: The Hub Card 52 receives Newbridge Cells from all of the UCS Cards 50 at a maximum rate of 200 Mbs and must buffer them in Queues before transmitting them on the UP-ISL 54 to the Switching Shelf 46. FIG. 5 shows the case of a `Standalone Hub` where the data going on the UP-ISL 54 is simply looped back (64) to the DN-ISL 56. Separate Queues must be maintained for the different levels of ATM Cell priority.

Detailed Description Text (20):

4. Output Queuing 70: Since the Drop Bus 68 is operating at 800 Mbs, the UCS Card 50 may receive more Newbridge Cells than it can instantaneously deal with. To prevent Cell loss an Output Queue 70 is required. UCS Cards must take into account the ATM Cell's priority when servicing the Output Queue so that different Qualities of Service can be provided.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L8: Entry 3 of 7

File: USPT

May 29, 2001

DOCUMENT-IDENTIFIER: US 6240066 B1

TITLE: Dynamic bandwidth and buffer management algorithm for multi-service ATM switches

Application Filing Date (1):
19980203

CLAIMS:

3. A method for dynamically allocating buffer and bandwidth resources as claimed in claim 1, wherein said computing step further comprises controlling the rate at which a selected one of said shared bandwidth resources and buffer resources are to be allocated until it is determined that the current bandwidth BW and buffer B resource allocation is biased towards a larger one of buffer and bandwidth allocation, respectively.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

L8: Entry 4 of 7

File: USPT

Sep 8, 1998

DOCUMENT-IDENTIFIER: US 5805203 A

TITLE: Time division, multiplexed, shared bandwidth communication system

Abstract Text (1):

A time division, multiplexed, shared bandwidth, high speed data system for use over CATV and broadcast equipment, which allows transmission rates to be adjusted on a per receiver basis in both forward (downstream) and return (upstream) directions, utilizes a time division multiplexed method to transmit information from a service provider to a plurality of users (downstream) using the equivalent of horizontal sync pulses to define the start and end of a period of transmission assigned to a specific user. The transmission is grouped into frames which typically contain five hundred to several thousand transmission periods. In each frame, at least one transmission period is reserved for each actively connected user which is used as a connection control channel. The remaining transmission periods are shared by all users and are dynamically allocated to transfer data. Hence, each user is allocated a small fixed bandwidth control channel and is provided with additional shared bandwidth for data transfers as needed. Transmission from the user to the service provider (upstream) is also time division multiplexed. Each user is allocated a period of time relative to the occurrence of the start of a frame, adjusted for the round trip propagation delay between the service provider and user, which is used as an upstream control channel which can be used to request additional shared upstream bandwidth for data transfers as needed.

Application Filing Date (1):

19960521

Brief Summary Text (5):

Among the obstacles which must be overcome in providing high speed digital services to broadcast and CATV users, are varying reception quality amongst the user base, adaptability to combination CATV systems in which the cable plant is a mixture of uni-directional and bi-directional cable, a noisy cable return channel spectrum (generally <42 MHz), varying length from user to bi-directional node or cable plant head end, and controlling the usage of shared bandwidth, both upstream and downstream, amongst multiple users.

Brief Summary Text (10):

The system utilizes a time division multiplexed method to transmit information from a service provider to a plurality of users (downstream), which has elements similar to that of conventional color television broadcast. This allows the use of standard television components in the customer equipment, substantially reducing the cost of the customer equipment. In this system the equivalent of horizontal sync pulses are used to define the start and end of a period of transmission assigned to a specific user. The transmission is grouped into frames which typically contain five hundred to several thousand transmission periods. In each frame, at least one transmission period is reserved for each actively connected user which is used as a connection control channel. The remaining transmission periods are shared by all users and are dynamically allocated to transfer data. Hence, each user is allocated a small fixed bandwidth control channel and is provided with additional shared bandwidth for data transfers as needed.

Detailed Description Text (3):

The system of the present invention provides a high speed digital service to users via CATV or broadcast, optionally using a telephone modem return path. Data is transmitted to and received from a plurality of users within discrete periods of time division multiplexed frames. There are two unique data channels provided to each user, a small fixed bandwidth control channel, and a dynamically allocated, shared bandwidth data channel. Each channel has an individually adjustable transmission method which is used to overcome noise. The control channel is used to manage the connection between the user and the service node including allocating bandwidth for the data channel and changing the transmission method. The data channel is used to relay information requests and responses between a user and information service provider.

[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L8: Entry 5 of 7

File: USPT

Oct 3, 1995

DOCUMENT-IDENTIFIER: US 5455826 A

TITLE: Method and apparatus for rate based flow control

Application Filing Date (1):

19940628

Brief Summary Text (15):

The disclosed flow control system includes a shared bandwidth pool on the network link that is shared among the multiple connections between the transmitting node and the receiving node. The system further includes a connection specific bandwidth allocation associated with each one of the multiple connections between the transmitting node and the receiving node. Each connection specific bandwidth allocation is available for transmitting DTUs over the associated connection. Thus, there is a connection specific bandwidth allocation reserved for each possible connection that is established between the transmitting node and the receiving node.

Detailed Description Text (15):

The flow control circuit reserves bandwidth sufficient to support a data rate of Min/T from the first network node 5 to the second network node 55 over a given VC. The remaining bandwidth is a shared bandwidth pool maintained by the flow control circuit, for transmitting data over any VC between the first network node 5 and the second network node 55. The memory 25 in the first network node 5 further contains a set of one or more data messages 27, the data messages 27 containing data to be transmitted by the first network node 5 to the second network node 55.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#)

[Generate Collection](#)

L8: Entry 7 of 7

File: DWPI

Jan 18, 2001

DERWENT-ACC-NO: 2001-299716

DERWENT-WEEK: 200382

COPYRIGHT 2005 DERWENT INFORMATION LTD

TITLE: Media access control layer for packet-centric wireless point to multipoint telecommunication system, has resource allocation device that allocates shared bandwidth among customer precise equipment stations

Basic Abstract Text (1):

NOVELTY - Media access control (MAC) layer has resource allocation device that allocates shared bandwidth among customer precise equipment (CPE) stations (294d,294e). Wireless base station (302) coupled to data network (142). CPE stations are coupled to base station in wireless communication over a shared bandwidth using a packet-centric protocol. Host computers (120d,122d) are coupled to CPE through LAN.

PF Application Date (4):

19980710

PF Application Date (5):

19990709

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L12: Entry 8 of 31

File: USPT

Jul 22, 2003

DOCUMENT-IDENTIFIER: US 6597689 B1
TITLE: SVC signaling system and method

Application Filing Date (1):
19981230

Brief Summary Text (23):

Despite these problems associated with blocking and oversubscription, most industry efforts have not been focused on prioritizing the incoming data into service requirements. As mentioned earlier, the various data received by a DSLAM may have specific service requirements (e.g., bandwidth, time synchronization) which may be affected by the blocking phenomenon. To overcome this problem, Quality of Service (QoS) could be implemented to "fairly" prioritize the various data received by the DSLAM so that some data can be serviced before other data, while at the same time, ensuring that all the data received are somehow "fairly" serviced. As known to those skilled in the art, an unexploited feature of ATM is the the ATM Forum's ATM service categories.

Brief Summary Text (30):

As mentioned above, a DSLAM that incorporates QoS to prioritize the various types of incoming data would be desirable. With such QoS support, a DSLAM would be better equipped to handle all the different service requirements of the incoming data.

Brief Summary Text (52):

The IMAS also supports end-to-end ATM over DSL. Because the IMAS is a fully compliant ATM switch that supports UNI 3.1, UNI 4.0, and PNNI 1.0 signaling, service providers can build end-to-end ATM networks over DSL from the subscriber's premise through a carrier ATM backbone network. With fully integrated call admission control (CAC) and multiple queues implementing weighted fair queueing, the telecommunications service provider can deploy high speed DSL bandwidth to various types of subscribers and guarantee Quality of Service. The IMAS concentrates and switches traffic from multiple subscribers utilizing ATM as the layer 2 protocol over the DSL access network to transparently transport layer 3 protocols such as IP and IPX.

Detailed Description Text (19):

In the LIU, some speed matching operation is necessary because DSL supports a lower bandwidth than the CSC and the telecommunications backbone. Once the data arrives into the LIUs in the DSLAMs, whether downstream from the telecommunications backbone or upstream from the subscribers via the DSL lines, the DSLAM handles the data based on a priority scheme, called priority queueing mechanism, in accordance with one embodiment of the present invention.

Detailed Description Text (20):

The priority queueing mechanism uses a combination of round robin and weighted fair queueing to resolve data access and bus contention issues by prioritizing which data will be handled and in what order. This prioritization is based on whether the data flow is downstream or upstream and whether the data is directed toward or away from the set of weighted queues in the line cards. Among subscribers, the priority

queueing mechanism ensures equal access to the DSLAM's communication resources; that is, the arbiter in the line card uses round robin to process data at all the ports, one port at a time. Once data has been processed in one port, the arbiter moves on to the next port to handle the data associated with this next port, and so on for all the ports.

Detailed Description Text (21):

For a given subscriber (or port), the priority queueing mechanism ensures some sort of fairness among the different types of data that are being accessed. Because each type of data may have different Quality of Service requirements (streaming video or audio v. static image), the DSLAM services each type of data in accordance with their respective service requirements. Some data will be serviced more frequently than others, but in accordance with the fairness aspect of the algorithm, all data will be serviced at some point.

Detailed Description Text (101):

Depending on the DSL technology supported, each line card may have a microprocessor-based computing environment inside. For those ADSL line cards using DMT, the line card has a DSP per modem for managing the modem. For those ADSL line cards using CAP, the line card has one processor for the entire line card for handling modem-specific housekeeping. For SDSL line cards, the line card contains a processor for debug and general modem-management purposes. Appropriate buffers are provided in each line card for priority queueing and other data management purposes.

Detailed Description Text (150):

The UTOPIA mux 253 is coupled to a filter 251 via line 275. This filter 251 determines whether the connection is valid or not. The filter 251 checks the validity of the connection based on an index to a look-up table 252 that corresponds to the VPI/VCI number of the ATM cell. The look-up table 252 is accessed via line 273. If the connection is not valid, the data is discarded. This look-up table has all the valid connection information, including translation functions, special marking, routing information, queue number and status information (e.g., discard state). To communicate with other devices in the line card, the filter 251 is coupled to the line card bus 272.

Detailed Description Text (151):

A buffer memory 247 is also provided. The primary purpose of this buffer memory 247 is to temporarily hold data that was delivered to it from the filter 251 and the framers 243 and 245. Conversely, the buffer memory 247 also holds data prior to its delivery to the filter 251 or the two framers 243 and 245. Any user traffic (i.e., ATM traffic) that arrives into the filter 251 via the UTOPIA mux 253 is held in the buffer memory 247 before it is delivered to any one of the destination DSL-side ports 235-242. Similarly, any user traffic that arrives into the framers (either framer 243 or framer 245) is held in the buffer memory 247 before it is delivered to the filter 251 on the CSC-side. The buffer memory 247 will be discussed in greater detail below in the PRIORITY QUEUEING MECHANISM subsection.

Detailed Description Text (152):

The data is organized according to port and priority. A buffer manager 248 is coupled to the buffer memory 247 via line 274. An arbiter 249 is coupled to the buffer manager 248 via line 276. The arbiter 249 resolves contention on the line card bus 272 via an arbitration scheme. Specifically, the arbiter 249 decides which of the many components that are coupled to the line card bus (i.e., filter 251 and framers 243 and 245) gets access to the bus for delivery of data into and out of the buffer memory 247 via a predetermined arbitration scheme. In accordance with one embodiment of the present invention, the arbitration scheme handles bus access (and hence, buffer memory access) according to the following priorities: (1) downstream data in the filter; (2) upstream data in the buffer memory; (3) upstream data in the framers; and (4) downstream data in the buffer memory. The PRIORITY

QUEUEING MECHANISM subsection below will provide more details on the arbitration scheme. In one embodiment of the present invention, this selection logic is implemented in hardware.

Detailed Description Text (162):

PRIORITY QUEUEING MECHANISM

Detailed Description Text (163):

The priority queueing mechanism will now be discussed. The priority queueing mechanism determines how and when data is transferred across the line card bus 272. In one embodiment, this priority queueing mechanism is implemented in hardware with the buffer memory 247, buffer manager 248, and arbiter 249. First, the various queues and FIFOs will be discussed. Second, the operation of the priority queueing mechanism will then be discussed with respect to the queues and FIFOs. In accordance with one embodiment of the present invention, the priority queueing mechanism handles ATM cells via a combination of round robin and weighted fair queueing (WFQ). Weighted fair queueing is discussed in detail in the WEIGHTED FAIR QUEUEING subsection below. So, when the patent specification mentions the transfer or delivery of data, this implies the transfer or delivery of cells.

Detailed Description Text (164):

PRIORITY QUEUEING MECHANISM--QUEUES AND FIFOs

Detailed Description Text (165):

The various components coupled to the bus 272 contain queues or first-in first-out (FIFO) buffers for handling all the data traffic. Some components have more and larger queues than other components.

Detailed Description Text (166):

As shown in FIG. 9, the buffer memory 247 includes a plurality of queues. The total number of queues can vary from one implementation to another. In this example, 132 queues are shown in groups of four, where each group of four represents a port, whether DSL-side or telecommunications backbone-side. If this line card had thirty-three ports (32 DSL-side and 1 CSC-side), 132 queues are used in this buffer memory 247. Thus, 128 queues are associated with downstream data, while 4 queues are associated with upstream data. If user data is temporarily stored in one of the downstream queues, it originally came from the telecommunications backbone via the chassis switch card in the IMAS and is ready to be transferred to the designated DSL-side port for delivery to the corresponding subscriber. Similarly, if user data is temporarily stored in one of the upstream queues, it originally came from a subscriber via the DSL line and framer (243 or 245) and is ready to be transferred to the telecommunications backbone via the UTOPIA mux 253 (see FIG. 8) in the line card and the chassis switch card.

Detailed Description Text (167):

To use another example, if eight DSL-side ports and one CSC-side port are provided in this line card, the buffer memory 247 must be large enough to accommodate nine total ports. For four queues per port, this particular line card uses at least thirty-six queues in the buffer memory 247. In one embodiment of the present invention, sixty-four queues are provided in which four queues are associated with a port. Thus, sixty-four queues can support sixteen ports.

Detailed Description Text (168):

Each queue in the buffer memory 247 is large enough to overcome any speed mismatch problems and provide enough buffering for bursty traffic. In one embodiment, each queue can hold anywhere from 128 cells to 2,048 cells. Although an ATM cell is 53 bytes long, the cell size used in the queues is 64 bytes. The actual size of the queues is programmable and can be adjusted to suit particular needs.

Detailed Description Text (169):

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L12: Entry 19 of 31

File: USPT

Sep 4, 2001

DOCUMENT-IDENTIFIER: US 6286052 B1

TITLE: Method and apparatus for identifying network data traffic flows and for applying quality of service treatments to the flows

Application Filing Date (1):
19981204

Brief Summary Text (9):

Computer networks include numerous services and resources for use in moving traffic throughout the network. For example, different network links, such as Fast Ethernet, Asynchronous Transfer Mode (ATM) channels, network tunnels, satellite links, etc., offer unique speed and bandwidth capabilities. Particular intermediate devices also include specific resources or services, such as number of priority queues, filter settings, availability of different queue selection strategies, congestion control algorithms, etc.

Brief Summary Text (10):

Individual frames or packets, moreover, can be marked so that intermediate devices may treat them in a predetermined manner. For example, the Institute of Electrical and Electronics Engineers (IEEE), in an appendix (802.1p) to the 802.1D bridge standard, describes additional information for the MAC header of Data Link Layer frames. FIG. 1 is a partial block diagram of a Data Link frame 100 which includes a MAC destination address (DA) field 102, a MAC source address (SA) field 104 and a data field 106. In accordance with the 802.1Q standard, a user priority field 108, among others, is inserted after the MAC SA field 104. The user priority field 108 may be loaded with a predetermined value (e.g., 0-7) that is associated with a particular treatment, such as background, best effort, excellent effort, etc. Network devices, upon examining the user priority field 108 of received Data Link frames 100, apply the corresponding treatment to the frames. For example, an intermediate device may have a plurality of transmission priority queues per port, and may assign frames to different queues of a destination port on the basis of the frame's user priority value.

Brief Summary Text (17):

A process executing at a given network entity, moreover, may generate hundreds if not thousands of traffic flows that are transmitted across the corresponding network every day. A traffic flow generally refers to a set of messages (frames and/or packets) that typically correspond to a particular task, transaction or operation (e.g., a print transaction) and may be identified by 5 network and transport layer parameters (e.g., source and destination IP addresses, source and destination TCP/UDP port numbers and transport protocol). Furthermore, the treatment that should be applied to these different traffic flows varies depending on the particular traffic flow at issue. For example, an on-line trading application may generate stock quote messages, stock transaction messages, transaction status messages, corporate financial information messages, print messages, data back-up messages, etc. A network administrator, moreover, may wish to have very different policies or service treatments applied to these various traffic flows. In particular, the network administrator may want a stock quote message to be given higher priority than a print transaction. Similarly, a \$1

million stock transaction message for a premium client should be assigned higher priority than a \$100 stock transaction message for a standard customer. Most intermediate network devices, however, lack the ability to distinguish among multiple traffic flows, especially those originating from the same host or server.

Brief Summary Text (24):

In another aspect of the invention, the local policy enforcer may establish a traffic flow state that includes the policy or service treatments specified by the policy server. It then monitors the traffic flows originating from the network entity looking for the given traffic flow. Once the given traffic flow is identified, the local policy enforcer applies the policy or service treatments set forth in the corresponding traffic flow state. For example, the policy enforcer may mark the packets or frames with a high priority DS codepoint. When the given traffic flow is complete, the application program may notify the flow declaration component, which, in turn, signals the end of the traffic flow to the local policy enforcer. The policy enforcer may request authorization from the policy server to release or otherwise discard the respective traffic flow state.

Detailed Description Text (7):

FIG. 3 is a partial block diagram of local policy enforcer 210. Local policy enforcer 210 includes a traffic flow state machine engine 310 for maintaining flow states corresponding to host/server 222 traffic flows, as described below. The traffic flow state machine engine 310 is coupled to a communication engine 312. The communication engine 312 is configured to formulate and exchange messages with the policy server 216 and the flow declaration component 226 at host/server 222. That is, communication engine 312 includes or has access to conventional circuitry for transmitting and receiving messages over the network 200. The traffic flow state machine engine 310 is also coupled to several traffic management resources and mechanisms. In particular, traffic flow state machine engine 310 is coupled to a packet/frame classifier 314, a traffic conditioner entity 316, a queue selector/mapping entity 318 and a scheduler 320. The traffic conditioner entity 316 includes several sub-components, including one or more metering entities 322, one or more marker entities 324, and one or more shaper/dropper entities 326. The queue selector/mapping entity 318 and scheduler 320 operate on the various queues established by local policy enforcer 210 for its ports and/or interfaces, such as queues 330a-330e corresponding to an interface 332.

Detailed Description Text (24):

The application-level parameters may encompass a whole range of information relating to different aspects of the traffic flow from the application program 224. For example, application-level parameters include such information as user name (e.g., John Smith), user department (e.g., engineering, accounting, marketing, etc.), application name (e.g., SAP R/3, PeopleSoft, etc.), application module (e.g., SAP R/3 accounting form, SAP R/3 order entry form, etc.), transaction type (e.g., print), sub-transaction type (e.g., print on HP Laser Jet Printer), transaction name (e.g., print monthly sales report), sub-transaction name (e.g., print monthly sales report on A4 paper), application state (e.g., normal mode, critical mode, primary mode, back-up mode, etc.). For a video streaming application, the application-level parameters might include user name, film name, film compression method, film priority, optimal bandwidth, etc. Similarly, for a voice over IP application, the application-level parameters may include calling party, called party, compression method, service level of calling party (e.g., gold, silver, bronze), etc. In addition, for World Wide Web (WWW) server-type applications, the application-level parameters may include Uniform Resource Locator (URL) (e.g., <http://www.altavista.com/cgi-in/> query?

pg=aq&kl=en&r=&search=Search&q=Speech+ne ar+recognition), front-end URL (e.g., <http://www.altavista.com>), back-end URL (e.g., query?

pg=aq&kl=en&r=&search=Search&q=Speech+near+recognition), mime type (e.g., text file, image file, language, etc.), file size, etc. Those skilled in the art will recognize that many other application-level parameters may be defined.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

L7: Entry 1 of 3

File: USPT

Sep 30, 2003

DOCUMENT-IDENTIFIER: US 6628615 B1
TITLE: Two level virtual channels

Application Filing Date (1):
20000118

Brief Summary Text (5):

Network connected multiprocessor systems typically comprise nodes which communicate through a switching network. The nodes may be uni-processor workstations or bus based shared memory multiprocessor systems (SMP). A node may also be an I/O subsystem such a disk array which itself contains an I/O processor. In such systems, a variety of traffic may be communicated over the inter-connection network: shared memory coherence and data messages, TCP/IP packets, disk blocks, user level messaging, etc.. Each type of traffic relies on certain properties of the network to provide the type of service the producers and consumers of that traffic expect. In some cases latency is critical, such as with shared memory coherence traffic and with some types of user level messages. In other cases, throughput is more critical than latency, such as with disk accesses. In some cases, a quality of service guarantee in terms of latency or bandwidth is required. The challenge for the interconnection network in such systems is to provide appropriate characteristics for each data type. Except for the quality of service case, this typically involves balancing requirements of one data type against those of another. Various types of inter-connection networks have been devised to address the general problem of providing good latency and throughput for a variety of traffic types. Most of these techniques have been developed in the context of packet switched (as opposed to circuit switched) networks. In these networks, the original message to be transmitted is decomposed into two smaller units. At one level, a message is broken into packets which may be fixed or variable in length. At the next level, packets are broken into fixed sized 'flits'. A flit is the fundamental data unit that can be interleaved at the lowest level of the network (i.e. the switching elements and the physical wires that connect them). The flit is also the level at which most techniques for enhancing network latency and throughput have been deployed.

Brief Summary Text (7):

A more recent development in packet switched networks for multi-processors is "virtual channels". Each physical channel, i.e., wire link between switching elements, is conceptually partitioned amongst multiple 'virtual' channels. Each virtual channel includes a physical virtual channel buffer on the switching element. The virtual channels are multiplexed across common physical channels, but otherwise operate independently. A blocked packet flit on one virtual channel does not block packet flits on a different virtual channel over a common physical channel.

Brief Summary Text (8):

Virtual channel networks provide better network utilization and reduce average communication latency by allowing data on one virtual channel (or lane) to overtake data on a different virtual channel when there is contention downstream on one channel but not on another. Another desirable property is guaranteed ordering of

transmissions on each channel and the ability to prioritize different data types. One factor that mitigates the improvement in network utilization is fragmentation of bandwidth on network links due to underutilized network flits. This can occur when data types assigned to different virtual channels are smaller than the flit. If these data types are communicated frequently, but not frequently enough to allow multiple of them to be packed into a flit, the flits become under utilized which can result in network under-utilization. Furthermore, there is a motivation to make flits large to increase the payload to overhead ratio, which only exacerbates the problem. Also, if the flit size is optimized for communication of large objects such as IP packets, the network may not be suitable for communication of smaller objects such as cache lines.

Brief Summary Text (15):

Further according to this scheme, first level channels are divided into two (2) channel classes: Latency Sensitive and Bandwidth Sensitive. Within each class each first level channel is assigned a unique priority. Flits on higher priority first level channels overtake flits on lower priority channels. Second level channels provide a dedicated connection between two system nodes. The end points of a second level channel may or may not reside on different system nodes. When two end points are connected, a second level channel is formed and assigned a globally unique second level channel id. First level channels flow control flits at the link level. Second level channels flow control packets at the network interface level (i.e. end to end). The transport agent breaks messages into packets, if necessary, and passes them to the second level virtual channels specified by the agent above the transport level (e.g., Non-Uniform Memory Access "NUMA" controller, session layer of a TCP/IP stack, etc.). If a latency sensitive message has a length $>M$, the transport agent rejects the request and returns an error condition code. Second level channels split bi-modal messages into two parts. The first N bytes of the message are passed to a first level latency sensitive channel and the remainder is sent to a first level bandwidth sensitive channel nearest in priority.

Brief Summary Text (16):

Advantageously, the system of the invention achieves greater network utilization in systems that require fine grain communication such as coherence controllers or, in systems that do course grain communications such as TCP/IP.

Detailed Description Text (6):

As further shown in FIG. 3, a second level channel 201 includes a frame buffer 307, 309 at both the source side 321 of the second level channel and at the destination side 322, respectively. The source side frame buffer 307 includes a number of slots, slots.sub.0, . . . , slots.sub.n that are filled in the order in which the buffer receives packet transfer requests from the transport agent. Packets are sent to the destination node in that same order and if the ordering is lost in passing through the network (e.g., adaptive routing) the order is recovered in the destination side frame buffer 309. This is accomplished through the use of sequence numbers transmitted with each packet which number is simply the index of the frame buffer slot corresponding to the packet. It is understood that recovery of lost packets may also be facilitated using this sequence number and a timeout mechanism. As shown in FIG. 3, each frame buffer slot includes four (4) pieces of information about packets that are in transit across the network: 1) a send side slot including a valid bit 302; an identifier of the message (message id) from which the packet came 303; 3) the base address in memory of the packet payload to be sent 304; and, 4) a transfer count 305 that is initialized with the payload size of the packet and then decremented as packet flits are passed to the network. The memory address of the packet payload may reference an on-chip location, as in the case of a Non-Uniform Memory Access (NUMA) coherence controller which generates request and response transactions in a collection of FIFO's. In that example, the memory address is actually the FIFO's out pointer.

Detailed Description Text (7):

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
[First Hit](#) [Fwd Refs](#)

[Generate Collection](#)

L13: Entry 12 of 23

File: USPT

Feb 6, 2001

DOCUMENT-IDENTIFIER: US 6185520 B1

TITLE: Method and system for bus switching data transfers

Application Filing Date (1):

19980522

Detailed Description Text (20):

It should be noted that in addition to the capabilities described above, switch 201 provides for "store and forward" data bursts where a data transaction from an initiator is buffered and forwarded at a later time (e.g., when the target device is not busy). Switch 201 also supports "cut-through" of data bursts, where successive accesses by an initiator are streamlined for efficiency (similar to the cut-through techniques of LAN switches). Additionally, it should be noted that switch 201 includes the computational resources which enable it to examine the contents of a data frame from a device and prioritize outgoing frames based upon user-defined criteria. This is similar to a multiple output queue technique supported in modern LAN switches (e.g., such as the CoreBuilder 3500 from 3Com Corp.).

Detailed Description Text (33):

In step 703, the switch of the present invention builds a prioritization table which indicates relative priority among various data transactions. For example, the prioritization table can be built such that the data transactions can be prioritized on a "data stream" basis (e.g., TCP socket, etc.), target device basis (shared bus), initiator device basis, or some combination of factors.

Detailed Description Text (37):

In step 707, the data transaction received from the first initiator is prioritized with respect to any other transactions received from other initiators. As described above, a priority selection criteria unit prioritizes among data transactions by selecting the appropriate output queue to be forwarded.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)